What AI job loss estimates actually measure

A source audit of exposure claims, workforce denominators and the evidence needed before projecting losses

Research by The Actual Factual | Brief 002 | Edition 1 | 14 November 2026

Purpose and scope

Reviewed 14 November 2026. This brief compares selected original studies and Census documentation; it is not a real-time count of AI layoffs or a forecast for a particular community. It identifies a reproducible route to demographic analysis without inventing population estimates.

Executive summary

The ILO's 2025 estimate places one in four workers worldwide in occupations with some generative AI exposure. It does not report that one in four jobs has disappeared. [1]

Exposure estimates, workplace productivity results and measured employment outcomes answer different questions. Studies should retain their geography, population, outcome and version when used in public arguments. [2] [3] [4] [5]

A Census population can supply a denominator, but it cannot turn a task-exposure score into a probability of dismissal. Demographic comparisons require matching occupation definitions, survey weights and uncertainty estimates. [6] [7] [8] [9]

Begin with the original claim

The ILO's May 2025 index combines task information, worker input, expert review and model-based scoring. It organizes occupations into exposure gradients; 3.3 percent of global employment is in its highest category. The authors judge transformation more likely than wholesale replacement because occupations contain tasks requiring human input. These are estimates of technological exposure, not observed redundancies. [1]

An IMF staff discussion note published in 2024 estimated that almost 40 percent of global employment was exposed to AI. Its framework separately considers whether AI might complement a job. Comparing that percentage with the ILO figure as if one were a revision of the other would obscure differences in the technology assessed and the exposure framework. Neither percentage should be relabelled a worldwide dismissal rate. [2]

A 2023 study by researchers associated with OpenAI, OpenResearch and the University of Pennsylvania estimated that roughly 80 percent of U.S. workers could have at least 10 percent of their tasks affected by GPT-related technologies. Its unit is potential task exposure. Eighty percent is not the share of jobs eliminated, and 10 percent of tasks is not necessarily 10 percent of paid working time. The institutional affiliations are disclosed because a technology provider participated in the research. [3]

Separate capability from outcomes

The published 2025 version of Generative AI at Work studied 5,172 customer-support agents. It found a 15 percent average increase in issues resolved per hour, with different effects by experience and skill. That is evidence about productivity in the studied setting. It does not measure a 15 percent economy-wide reduction in staff. Earlier working-paper versions used different figures; quote the version actually consulted. [4]

Humlum and Vestergaard linked Danish chatbot-adoption surveys to administrative labor records. The March 2026 working-paper abstract reports no detectable average earnings or recorded-hours effects, ruling out effects larger than 2 percent over its two-year post-ChatGPT window, alongside changes in tasks. This is evidence against large early average effects in that study, not proof that no individual was affected or that future losses cannot occur. [5]

Our interpretation: an oversight brief should track at least four stages separately - technical capability, actual adoption, changes in work, and labor outcomes. A result at one stage does not establish the next. Before attributing workforce reductions to AI, request the deployment dates, affected functions, headcount and hours series, hiring changes, outsourcing decisions and the employer's explanation. Test plausible alternative explanations rather than choosing the most compelling headline.

Choose a population that matches the question

For a U.S. occupation-based inquiry, the 2024 ACS one-year detailed table B24114 defines its universe as the civilian employed population aged 16 and over. Its metadata identifies graphic designers under estimate B24114_171E and margin-of-error variable B24114_171M. That is a concrete starting point for a national occupation count; it is not a count of everyone who has design skills or every person employed by a design business. This brief does not report an extracted graphic-designer estimate. [6]

For a question such as how many women in a particular occupation could face changed tasks, first specify geography, survey year, employment status, occupation, and the demographic variable. Use the relevant published table where available; otherwise document the PUMS selection and person weights. The Census PUMS documentation supplies variable definitions, code lists and weight guidance. Do not count sample records as if each represented one resident. [7]

Occupation codes need a documented crosswalk. Census explains that classification changes can move workers between categories and limit comparisons over time. Matching labels by eye is insufficient when a study uses a different occupational system. Publish unmatched and split categories and explain allocation assumptions. Otherwise, a precise-looking total may conceal an uncertain mapping. [8]

Show the calculation and its limits

Illustration only: suppose a clearly defined population contains 10,000 workers and an explicitly assumed scenario classifies 20 percent as potentially exposed. The multiplication is 10,000 × 0.20 = 2,000 workers in that scenario. Neither input is an actual Census estimate here. The result is not 2,000 predicted layoffs, and it has no forecast year unless a separate model supplies one.

If a published study instead says that 20 percent of tasks within an occupation meet its exposure threshold, multiplying that percentage by workers changes the unit of measurement. The result cannot be described as a

count of exposed people without an additional, justified definition. Publish the numerator, denominator, unit, geography, period and assumptions beside every calculation.

Survey uncertainty and scenario uncertainty are different. Census publishes 90 percent margins of error and provides instructions for statistical comparisons. Those margins address sampling variability; they do not validate an assumed AI effect. A hypothetical exposure percentage must remain labelled even when the underlying worker estimate is official. [9]

Do not manufacture demographic effects

A group's share of an occupation can explain how a common scenario distributes numerically across that occupation. It does not establish that the technology affects every subgroup at the same rate. Conversely, applying different rates by sex, race, disability or age requires an evidentiary basis or an explicit scenario assumption. Keep occupational composition, within-occupation exposure and observed outcomes separate.

Recommended reporting rule: show both the estimated number and the percentage of each group's relevant workforce. State whether categories overlap, explain missing or suppressed cells, and assess whether apparent differences exceed sampling uncertainty. Census statistical guidance supports testing comparisons rather than ranking estimates by size alone. [9] Without these steps, a demographic headline can confuse a larger population with a higher individual risk.

Evidence record

The table identifies the record, the finding it supports, and the limit of that evidence. Numbered citations open the original source.

Date or record	Evidence and limits
17 March 2023	**U.S. task exposure.** GPT research describes potential effects on tasks across the U.S. workforce. [3] **Limit:** Capability rubric, not a count of dismissals.
January 2024	**Global exposure and complementarity.** IMF staff distinguish exposure from whether technology may complement work. [2] **Limit:** The framework is not interchangeable with an ILO generative AI index.
4 February 2025	**Observed workplace productivity.** The published customer-support study measures issues resolved per hour. [4] **Limit:** One setting cannot establish economy-wide employment effects.
20 May 2025	**Global generative AI exposure.** ILO publishes a revised occupational exposure framework. [1] **Limit:** Exposure categories do not supply a dismissal probability.
March 2026 revision	**Early labor outcomes in Denmark.** The working paper measures earnings and hours using linked survey and administrative records. [5] **Limit:** Country, time window and estimated averages constrain interpretation.

Date or record	Evidence and limits
2024 survey vintage	**U.S. Census denominator.** ACS B24114 defines a civilian employed population aged 16 and over and includes detailed occupations. [6] **Limit:** Population estimates do not measure AI causation.

Claims findings and unknowns

Claim 1

Claim under examination: A global exposure percentage tells us how many local jobs will vanish.

Finding: An exposure estimate describes its defined population and threshold. [1] [2]

What remains unknown: Local adoption, task mix, employer responses and future employment outcomes.

Claim 2

Claim under examination: Higher productivity means the same percentage fewer workers.

Finding: The cited workplace study measures output per hour. [4]

What remains unknown: Whether demand, hours, staffing, pay or service quality change afterward.

Claim 3

Claim under examination: An official Census denominator makes a scenario a forecast.

Finding: Official population data and an assumed effect remain separate inputs. [6] [7] [9]

What remains unknown: The validity of the effect assumption and its transfer to the selected population.

Questions for decision makers

These are recommended oversight questions. They identify records that could support an answer and do not presume a particular conclusion.

Question 1

What precisely does the headline percentage measure?

Why this matters: A task share, a worker share and a job-loss rate are different quantities. [1] [2] [3]

Who can answer: Study authors, analysts and communications teams

Records to request: Original study, version, threshold definition, numerator and denominator.

Question 2

Which population was used for the local calculation?

Why this matters: An estimate must match the geography, occupation, year and employment status in the question. [6] [7]

Who can answer: Workforce agency or commissioning organization

Records to request: Census table and cells or PUMS query, filters, weights and extraction date.

Question 3

How were occupation codes connected between datasets?

Why this matters: A crosswalk can change who is counted and conceal ambiguous matches. [8]

Who can answer: Research team and data contractor

Records to request: Code versions, crosswalk, split-category allocations and unmatched records.

Question 4

What evidence connects deployment to the reported labor outcome?

Why this matters: An outcome following an announcement is insufficient to establish its cause. [4] [5]

Who can answer: Employer leadership, worker representatives and independent researchers

Records to request: Deployment chronology, comparison group, headcount, hiring, hours, wages, separations and outsourcing records.

Question 5

Are demographic differences measured or assumed?

Why this matters: Group counts, within-group rates and differences in treatment should not be conflated. [7] [9]

Who can answer: Researchers and public officials using the estimates

Records to request: Subgroup definitions, unweighted sample sizes, weights, margins of error, statistical tests and scenario assumptions.

Connections and their limits

These relationships are a reading aid. A connection does not establish causation, strength of influence, or motive. The interactive version is available on the report landing page.

Task exposure and Occupational classification

Documented relationship: Exposure research classifies work through occupational and task frameworks. [1] [3]

Limit: A category is not an individual dismissal forecast.

Occupational classification and Census population

Editorial analysis: A defensible population estimate requires a documented match between study codes and Census categories. [6] [8]

Limit: Ambiguous mappings must remain visible.

Census population and Demographic comparison

Documented relationship: PUMS documentation supplies variables and weighting guidance for population analysis. [7] [9]

Limit: This brief has not estimated demographic exposure counts.

AI adoption and Productivity

Documented relationship: The customer-support study evaluates productivity after an AI assistant was introduced. [4]

Limit: Its outcome is not the number of eliminated jobs.

Productivity and Employment outcomes

Open question: Do changes in output translate into staffing, pay or hours changes, and for whom? [4] [5]

Limit: This relationship must be investigated; it is not established by a single percentage.

Method and corrections history

Edition 1, reviewed 14 November 2026. Selected original research and official statistical documentation were checked for population, unit, outcome and version. This is a source audit and calculation guide, not an exhaustive literature review or independent replication of the studies. The numerical example is hypothetical. No local or demographic job-loss count is claimed. Analysis and recommended questions are labelled; future revisions should identify altered sources, study versions and calculations. Corrections history: initial edition, no prior corrections.

14 November 2026 | Edition 1. Initial publication of this research brief. No corrections recorded in this edition. Future substantive updates should identify what changed and the supporting evidence.

Web report and future updates

Source register

Each numbered citation and source title links to the original record. Source dates identify the cited version, not necessarily its current legal or operational status.

[1] Generative AI and Jobs A Refined Global Index of Occupational Exposure

International Labour Organization | 20 May 2025

Working Paper 140. Original institutional research; occupational exposure, not observed job losses.

[2] Gen AI Artificial Intelligence and the Future of Work

International Monetary Fund | January 2024

Staff Discussion Note 2024/001. Exposure and complementarity framework; staff analysis rather than an employment forecast.

[3] GPTs are GPTs An early look at the labor market impact potential of large language models

OpenAI | 17 March 2023

Original research summary and paper link. Provider-affiliated authorship disclosed; prospective task exposure.

[4] Generative AI at Work

The Quarterly Journal of Economics | 4 February 2025

Brynjolfsson, Li and Raymond. Published version used; differs from earlier working-paper sample and estimates.

[5] Still Waters Rapid Currents Early Labor Market Transformation under Generative AI

National Bureau of Economic Research | March 2026 revision

Humlum and Vestergaard, Working Paper 33777. Danish early-outcomes study; revised from Large Language Models Small Labor Market Effects.

[6] 2024 ACS B24114 variable definitions

U.S. Census Bureau | 2024 survey vintage

Detailed occupations for civilian employed people aged 16 and over. Graphic designer estimate and margin-of-error identifiers verified; no count extracted for this brief.

[7] American Community Survey PUMS Documentation

U.S. Census Bureau | 2024 documentation selected

Variable definitions, user guidance, weights and accuracy documentation organized by release year.

[8] About Occupation

U.S. Census Bureau | Reviewed 14 November 2026

Explains occupational coding and limits to comparisons across changing classifications.

[9] Instructions for Applying Statistical Testing to ACS Data

U.S. Census Bureau | 2024 instructions

Methods for margins of error and statistical comparisons. Survey uncertainty is distinct from AI-scenario uncertainty.